EFFICIENT SUM OF SPARSE OUTER PRODUCTS DICTIONARY LEARNING (SOUP-DIL) *

Saiprasad Ravishankar, Raj Rao Nadakuditi & Jeffrey A. Fessler
Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI 48109, USA
{ravisha, rajnrao, fessler}@umich.edu

arXiv:1511.06333v3 [cs.LG] 7 Jan 2016

ABSTRACT

The sparsity of natural signals in a transform domain or dictionary has been extensively exploited in several applications. More recently, the data-driven adaptation of synthesis dictionaries has shown promise in many applications compared to fixed or analytical dictionaries. However, dictionary learning problems are typically non-convex and NP-hard, and the usual alternating minimization approaches for these problems are often computationally expensive, with the computations dominated by the NP-hard synthesis sparse coding step. In this work, we investigate an efficient method for ℓ_0 "norm"-based dictionary learning by first approximating the training data set with a sum of sparse rank-one matrices and then using a block coordinate descent approach to estimate the rank-one terms. The proposed algorithm involves efficient closed-form solutions. In particular, the sparse coding step involves a simple form of thresholding. We provide a convergence analysis for the proposed block coordinate descent method. Our experiments show the promising performance and significant speed-ups provided by our method over the classical K-SVD scheme in sparse signal representation and image denoising.

1 INTRODUCTION

The sparsity of natural signals and images in a transform domain or dictionary has been extensively exploited in several applications such as compression, denoising, compressed sensing and other inverse problems. More recently, the data-driven adaptation of synthesis dictionaries called dictionary learning, has shown promise in applications (Elad & Aharon, 2006; Mairal et al., 2009; Ravishankar & Bresler, 2011) compared to analytical dictionaries such as wavelets or DCT.

Given a collection of training signals $\{y_i\}_{i=1}^N$ that are represented as columns of the matrix $Y \in \mathbb{R}^{n \times N}$, the dictionary learning problem is often formulated as follows (Aharon et al., 2006)

$$\text{(P0)} \ \min_{D,X} \|Y - DX\|_F^2 \ \text{ s.t. } \ \|x_i\|_0 \leq s \ \forall\, i, \ \|d_j\|_2 = 1 \, \forall\, j.$$

Here, d_j and x_i denote the columns of the dictionary $D \in \mathbb{R}^{n \times J}$ and sparse code matrix $X \in \mathbb{R}^{J \times N}$, respectively, and s denotes the maximum sparsity level (non-zeros in representations x_i) allowed for each training signal. The ℓ_0 "norm" counts the number of non-zero entries in a vector. The columns of the dictionary are constrained to have unit norm in (P0) to avoid the scaling ambiguity (Gribonval & Schnass, 2010). Variants of Problem (P0) include replacing the ℓ_0 "norm" for sparsity with an ℓ_1 norm or an alternative sparsity criterion, or enforcing additional properties (e.g., incoherence) for the dictionary D, or solving an online version (where the dictionary is updated sequentially as new training signals arrive) of the problem (Mairal et al., 2010).

Algorithms for (P0) or its variants (Engan et al., 1999; Aharon et al., 2006; Yaghoobi et al., 2009; Skretting & Engan, 2010; Mairal et al., 2010; Smith & Elad, 2013; Sadeghi et al., 2013) typically alternate[1] in some form between a *sparse coding step* (updating X), and a *dictionary update step*

*An extended version of this work is available online (Ravishankar et al., 2015) that contains greater details and the proofs of various theoretical convergence results stated in this work.

[1] Alternating algorithms belong to the class of block coordinate descent algorithms.

(solving for D). Some of these algorithms also partially update X in the dictionary update step. A few recent methods attempt to solve for D and X jointly in an iterative fashion (Rakotomamonjy, 2013). The K-SVD method (Aharon et al., 2006) has been particularly popular and demonstrated to be useful in numerous applications (Elad & Aharon, 2006; Ravishankar & Bresler, 2011). However, Problem (P0) is highly non-convex and NP-hard, and most dictionary learning approaches lack proven convergence guarantees. Moreover, the algorithms for (P0) tend to be computationally expensive (particularly alternating-type algorithms), with the computations usually dominated by the synthesis sparse coding step.

Some recent works (Spielman et al., 2012; Arora et al., 2013; Xu & Yin, 2013; Bao et al., 2014; Agarwal et al., 2014) have studied the convergence of (specific) synthesis dictionary learning algorithms. However, these dictionary learning methods have not been demonstrated to be useful in applications such as image denoising. Bao et al. (2014) find that their method, although a fast proximal scheme, denoises less effectively (typically $0.1 - 0.3$ dB worse) than the K-SVD method. Many prior works use restrictive assumptions (e.g., noiseless data, etc.) for their convergence results.

In this work, we focus on synthesis dictionary learning and investigate a learning problem with an ℓ_0 sparsity penalty. Our approach first models the training data set as an approximate sum of sparse rank-one matrices (or outer products). Then, we use a highly efficient and exact block coordinate descent approach to estimate the factors of the various rank-one matrices. We provide a convergence analysis of the proposed block coordinate descent method. Our numerical experiments demonstrate the promising performance and significant speed-ups provided by our method over the classicial K-SVD dictionary learning scheme in sparse signal representation and image denoising. Importantly, our algorithm is similar to just the dictionary update step of K-SVD, and yet performs comparably or much better and faster than K-SVD in applications.

Our work shares similarities with a recent dictionary learning approach (Sadeghi et al., 2014) that exploits a sum of outer products model for the training data. However, unlike the prior work that considered an ℓ_1 norm-based algorithm as a heuristic variation of K-SVD, we focus in this work on the harder ℓ_0 "norm"-based dictionary learning problem (with additionally required ℓ_∞ norm constraints), and provide a rigorous derivation of an efficient block coordinate descent algorithm for this case. We also briefly discuss various alternative scenarios/formulations, where the sum of outer products optimization approach is beneficial. Finally, unlike prior work (Sadeghi et al., 2014), we provide a detailed convergence analysis of the proposed algorithm for a highly non-convex problem, and moreover demonstrate its usefulness (both in terms of speed and quality of results) in sparse representation of natural signals and image denoising. We compare our method to other ℓ_0 "norm"-based methods in this work. A detailed investigation and comparison between various alternative formulations (including ℓ_1 formulations) is left for future work.

2 Problem Formulation and Algorithm

In this section, we discuss a dictionary learning problem formulation with an ℓ_0 penalty, and present an efficient algorithm for this non-convex problem, and discuss the algorithms' computational benefits. We also briefly discuss possible alternatives and extensions.

2.1 ℓ_0 Penalized Formulation

We consider a sparsity penalized variant of Problem (P0) (Bao et al., 2014) in this work. Specifically, we replace the sparsity constraints in (P0) with an ℓ_0 penalty $\sum_{i=1}^{N} \|x_i\|_0$. Next, with $C \triangleq X^T$, we express the matrix $DX = DC^T$ as a sum of (sparse) rank-one matrices or outer products $\sum_{j=1}^{J} d_j c_j^T$, where c_j is the jth column of C. This Sum of OUter Products (SOUP) then provides an approximation (not necessarily an exact decomposition) of the training data Y. The SOUP representation of Y is natural because it separates out the contributions of the various atoms in representing the data. It may also provide a natural way to set the number of atoms (degrees of freedom) in the dictionary. In particular, atoms of a dictionary whose contributions to the data (Y) representation error or modeling error are small could be dropped. Such a SOUP representation has been exploited before (Aharon et al., 2006; Smith & Elad, 2013). With the above modifications, we

have the following dictionary learning problem, where $\lambda^2 > 0$ is a weight:

$$\text{(P1)} \ \min_{\{d_j, c_j\}} \left\| Y - \sum_{j=1}^{J} d_j c_j^T \right\|_F^2 + \lambda^2 \sum_{j=1}^{J} \|c_j\|_0 \ \text{ s.t. } \ \|d_j\|_2 = 1, \ \|c_j\|_\infty \leq L \ \forall \, j.$$

As in Problem (P0), the matrix $d_j c_j^T$ in (P1) is invariant to joint scaling of d_j and c_j as αd_j and $(1/\alpha)c_j$, for $\alpha \neq 0$. The constraint $\|d_j\|_2 = 1$ helps remove this scaling ambiguity. We also enforce the constraint $\|c_j\|_\infty \leq L$, with $L > 0$, in (P1). Consider a dictionary D that has a column d_j that repeats. Then, in this case, the outer product expansion of Y in (P1) could have both the terms $d_j c_j^T$ and $-d_j c_j^T$ with c_j that is highly sparse, and the objective would be invariant to (arbitrarily) large scalings of c_j (i.e., non-coercive objective). The ℓ_∞ constraints on the columns of C (that constrain the magnitudes of entries of C) prevent such a scenario.

Problem (P1) is designed to learn the factors $\{d_j\}_{j=1}^J$ and $\{c_j\}_{j=1}^J$ that enable the best SOUP sparse representation of Y. (P1) is harder than a penalized sparse coding problem (with fixed dictionary). In particular, (P1) (and (P0)) remains non-convex even when the ℓ_0 "norm" is relaxed to an ℓ_1 norm.

Formulation (P1) is by itself useful for adaptive sparse representation (and eventual compression) of data. Learning sparse approximations of the data can also be an effective way to denoise the data. These applications are considered in Section 4. (P1) can also be directly used as a regularizer in inverse problems such as blind compressed sensing (Ravishankar & Bresler, 2015), where it can enable data-adaptive image reconstructions. In cases, when the dictionary adapted using (P1) is used for approximating a different group of test signals, any of the usual sparse coding techniques (Pati et al., 1993; Chen et al., 1998; Needell & Tropp, 2009) can be adopted to encode signals using a known dictionary. For example, in the experiments in Section 4, learned dictionaries are used to sparse code data using the Orthogonal Matching Pursuit (OMP) algorithm.

2.2 Alternative Formulations

Several variants of the SOUP learning Problem (P1) could be constructed, with interesting effects. For example, the ℓ_0 "norm" for sparsity could be replaced by the ℓ_1 norm (Sadeghi et al., 2014). Problem (P1) can also be modified to learn a doubly sparse dictionary (Rubinstein et al., 2010) by setting $d_j = \Phi a_j \ \forall \ j$ for a fixed orthonormal Φ, and employing a sparsity penalty for the a_j's. Another interesting alternative to (P1) involves enforcing p-block-orthogonality constraints on the dictionary D. The dictionary in this case is split into blocks (instead of individual atoms), each of which has p atoms that are orthogonal to each other. For $p = 2$, the constraints take the form $d_{2j-1}^T d_{2j} = 0$, $1 \leq j \leq J/2$. In the extreme (more constrained) case of $p = n$, the dictionary would be made of several (square) orthonormal blocks. For tensor-type data, Problem (P1) can be modified by enforcing the dictionary atoms to be in the form of a Kronecker product. The algorithm that we propose in Section 2.3 can be easily extended to accomodate such variants of Problem (P1). We do not fully explore such alternatives in this work due to space constraints, and a more detailed investigation, analysis, and comparison of these variants will be presented elsewhere.

2.3 Algorithm

We propose a block coordinate descent method to estimate the unknowns in Problem (P1). For each $1 \leq j \leq J$, we first solve (P1) with respect to c_j keeping all the other variables fixed (called *sparse coding step*). Once c_j is updated, we solve (P1) with respect to d_j keeping all other variables fixed (called *dictionary atom update step* or *dictionary update step*).

2.3.1 Sparse Coding Step

Minimizing (P1) with respect to c_j leads to the following non-convex problem, where $E_j \triangleq Y - \sum_{k \neq j} d_k c_k^T$ is a fixed matrix based on the most recent values of all other atoms and coefficients:

$$\min_{c_j} \ \left\| E_j - d_j c_j^T \right\|_F^2 + \lambda^2 \|c_j\|_0 \ \text{ s.t. } \ \|c_j\|_\infty \leq L. \tag{1}$$

The following proposition (see proof in Appendix A) provides the solution to Problem (1), where the hard-thresholding operator $H_\lambda(\cdot)$ is defined as

$$(H_\lambda(b))_i = \begin{cases} 0, & |b_i| < \lambda \\ b_i, & |b_i| \geq \lambda \end{cases} \tag{2}$$

with $b \in \mathbb{R}^N$, and the subscript i indexes vector entries. We assume that the bound $L > \lambda$ and let 1_N denote a vector of ones of length N. The operation "$\odot$" denotes element-wise multiplication, $\text{sign}(\cdot)$ computes the signs of the elements of a vector, and $z = \min(a, u)$ for vectors $a, u \in \mathbb{R}^N$ denotes the element-wise minimum operation, i.e., $z_i = \min(a_i, b_i)$, $1 \leq i \leq N$.

Proposition 1 *Given $E_j \in \mathbb{R}^{n \times N}$ and $d_j \in \mathbb{R}^n$, and assuming $L > \lambda$, a global minimizer of the sparse coding problem* (1) *is*

$$\hat{c}_j = \min\left(\left|H_\lambda\left(E_j^T d_j\right)\right|, L1_N\right) \odot \text{sign}\left(H_\lambda\left(E_j^T d_j\right)\right) \tag{3}$$

The solution is unique if and only if the vector $E_j^T d_j$ has no entry with a magnitude of λ.

2.3.2 Dictionary Atom Update Step

Minimizing (P1) with respect to d_j leads to the following non-convex problem:

$$\min_{d_j} \left\|E_j - d_j c_j^T\right\|_F^2 \quad \text{s.t.} \quad \|d_j\|_2 = 1 \tag{4}$$

Proposition 2 (see proof in Appendix B) provides the closed-form solution for (4).

Proposition 2 *Given $E_j \in \mathbb{R}^{n \times N}$ and $c_j \in \mathbb{R}^N$, a global minimizer of the dictionary atom update problem* (4) *is*

$$\hat{d}_j = \begin{cases} \frac{E_j c_j}{\|E_j c_j\|_2}, & \text{if } c_j \neq 0 \\ v_1, & \text{if } c_j = 0 \end{cases} \tag{5}$$

where v_1 is the first column of the $n \times n$ identity matrix. The solution is unique if and only if $c_j \neq 0$.

2.3.3 Overall Algorithm and Computational Properties

Fig. 1 shows the Sum of OUter Products DIctionary Learning (SOUP-DIL) Algorithm 1 for Problem (P1). The algorithm assumes that an initial estimate $\left\{d_j^0, c_j^0\right\}_{j=1}^J$ for the variables is provided. For example, the initial sparse coefficients can be set to zero, and the initial dictionary can be a known analytical dictionary such as the overcomplete DCT (Elad & Aharon, 2006). When $c_j^t = 0$, the setting $d_j^t = v_1$ in (9) in Algorithm 1 could also be replaced with other (equivalent) settings such as $d_j^t = d_j^{t-1}$ or setting d_j^t to a random unit norm vector. All of these settings have been observed to work well in practice. A random ordering of the atom/sparse coefficient updates in Fig. 1 (i.e., random j sequence) can also help in practice (in accelerating convergence) compared to cycling in the same order 1 through J every iteration. One could also alternate several times between the sparse coding and dictionary atom update steps for each j in Fig. 1. However, this would increase computation.

Note that the way in which block coordinate descent is performed for (P1) is important in order to obtain efficient closed-form solutions for the resulting sub-problems. For example, one could attempt to alternate between a block coordinate descent-based sparse coding step (for updating all coefficients in C) and a full dictionary D update step for (P1). However, because of the hard constraints on D, there does not exist a closed-form solution with respect to the entire dictionary, for fixed sparse codes. But, as we have shown, efficient closed-form solutions are indeed obtained when, for example, optimizing for D (or C) one column at a time in (P1). Various ways of achieving and ordering closed-form updates may be possible. However, we have proposed a particular highly efficient approach in Fig. 1 relying on the SOUP representation. Note that we do not update all the c_k's to maximum possible accuracy (using iterative block coordinate descent), every time a column of D is updated in Fig. 1. This would substantially increase the computation.

We now discuss the computational advantages of Algorithm 1. Define the sparsity factor α of the sparse matrix $C \in \mathbb{R}^{N \times J}$ as the non-negative number satisfying $\sum_{j=1}^J \|c_j\|_0 = \alpha N n$. This

SOUP-DIL Algorithm 1

Inputs : Training data $Y \in \mathbb{R}^{n \times N}$, weight λ, upper bound L, and number of iterations K.

Outputs : Columns $\left\{d_j^K\right\}_{j=1}^J$ of the learned dictionary, and the learned $\left\{c_j^K\right\}_{j=1}^J$.

Initial Estimates: $\left\{d_j^0, c_j^0\right\}_{j=1}^J$.

For $t = 1 : K$ **Repeat**

For $j = 1 : J$ **Repeat**

1. $C = \left[c_1^t, ..., c_{j-1}^t, c_j^{t-1}, ..., c_J^{t-1}\right]$, $D = \left[d_1^t, ..., d_{j-1}^t, d_j^{t-1}, ..., d_J^{t-1}\right]$.

2. **Sparse coding:**

$$b^t = Y^T d_j^{t-1} - CD^T d_j^{t-1} + c_j^{t-1} \tag{6}$$

$$c_j^t = \min\left(\left|H_\lambda\left(b^t\right)\right|, L1_N\right) \odot \operatorname{sign}\left(H_\lambda\left(b^t\right)\right) \tag{7}$$

3. **Dictionary atom update:**

$$h^t = Yc_j^t - DC^T c_j^t + d_j^{t-1}\left(c_j^{t-1}\right)^T c_j^t \tag{8}$$

$$d_j^t = \begin{cases} \frac{h^t}{\|h^t\|_2}, & \text{if } c_j^t \neq 0 \\ v_1, & \text{if } c_j^t = 0 \end{cases} \tag{9}$$

End

End

Figure 1: SOUP-DIL Algorithm 1 for Problem (P1). Superscript of t denotes the iterates in the algorithm. Both b^t and h^t above can be computed very efficiently via sparse operations.

translates to an average of αn non-zeros per row of C or $\alpha Nn/J$ non-zeros per column of C. It is easy to show (Ravishankar et al., 2015) that the computational cost of the $J \geq n$ inner iterations (sparse coding and atom update) in iteration t in Fig. 1 is dominated (for $N \gg J, n$) by $NJn + 2\alpha_m NJn + \beta Nn^2$, where α_m is the maximum sparsity factor of the estimated C's during the inner iterations, and β is the sparsity factor of the estimated C at the end of iteration t. Thus, the cost per iteration of the block coordinate descent SOUP-DIL Algorithm 1 is about $(1 + \alpha')NJn$, with $\alpha' \ll 1$ typically. On the other hand, the proximal alternating algorithm (that involves more parameters than our scheme) proposed very recently by Bao et al. (2015; 2014) has a per-iteration computational cost of at least $2NJn + 6\alpha NJn + 4\alpha Nn^2$. This is clearly more computation than SOUP-DIL. Assuming $J \propto n$, the cost per iteration of Algorithm 1 scales as $O(Nn^2)$. This is lower than the per-iteration cost of learning an $n \times J$ synthesis dictionary D using K-SVD (Aharon et al., 2006), which scales (assuming that the synthesis sparsity level $s \propto n$ and $J \propto n$ in K-SVD)[2] as $O(Nn^3)$.

As illustrated in Section 4.1, our algorithms converge in few iterations in practice. Therefore, the per-iteration computational advantages also translate to net computational advantages in practice. The low computational cost of our approach could be particularly useful for big data applications, or higher dimensional (3D or 4D) applications.

The computations in Algorithm 1 can also be easily parallelized. In particular, the matrix-vector products and truncated thresholding operations in the J ($J \ll N$ typically) inner iterations of each iteration t in Fig. 1 can be implemented efficiently using parallel computing (Grama et al., 2003).

3 Convergence Analysis

This section presents a convergence analysis for the Algorithm 1 for Problem (P1). The proposed algorithm is an exact block coordinate descent procedure for Problem (P1). However, due to the high degree of non-convexity involved, standard results on convergence of block coordinate descent methods (e.g., (Tseng, 2001)) do not apply here. We present some definitions and notations below, before stating our convergence results for Algorithm 1.

[2]When $s \propto n$ and $J \propto n$, the per-iteration computational cost of the efficient implementation of K-SVD (Rubinstein et al., 2008) also scales similarly as $O(Nn^3)$.

First, a necessary condition for $x \in \mathbb{R}^p$ to be a minimizer of a (proper) function $g : \mathbb{R}^p \mapsto (-\infty, +\infty]$ is that x is a *critical point* of g, i.e., $0 \in \partial g(x)$, where $\partial g(x)$ is the sub-differential of g at x (Rockafellar & Wets, 1997; Mordukhovich, 2006). Critical points are considered to be "generalized stationary points" (Rockafellar & Wets, 1997).

The constraints $\|d_j\|_2 = 1$, $1 \leq j \leq J$, in (P1) can instead be added as penalties in the cost by using barrier functions $\chi(d_j)$ (taking the value $+\infty$ when the norm constraint is violated, and zero otherwise). The constraints $\|c_j\|_\infty \leq L$, $1 \leq j \leq J$, can also be similarly replaced with barrier penalties $\psi(c_j)$. Then, (P1) can be written in an unconstrained form with objective

$$f(C,D) = f\left(c_1, ..., c_J, d_1, ..., d_J\right) = \left\|Y - \sum_{j=1}^{J} d_j c_j^T\right\|_F^2 + \sum_{j=1}^{J} \left\{\lambda^2 \|c_j\|_0 + \chi(d_j) + \psi(c_j)\right\}$$

For Algorithm 1, the iterates in the tth outer iteration are denoted as $\left\{c_j^t, d_j^t\right\}_{j=1}^J$ or as (C^t, D^t).

3.1 Main Results

Assume that the initial (C^0, D^0) satisfies the constraints in (P1). We then have the following result. (Due to space constraints, the proofs of results are presented elsewhere (Ravishankar et al., 2015).)

Theorem 1 *Let $\{C^t, D^t\}$ denote the iterate sequence generated by Algorithm 1 with training data $Y \in \mathbb{R}^{n \times N}$ and initial (C^0, D^0). Then, the objective sequence $\{f^t\}$ with $f^t \triangleq f\left(C^t, D^t\right)$ is monotone decreasing, and converges to a finite value, say f^*. Moreover, the iterate sequence $\{C^t, D^t\}$ is bounded, and all its accumulation points are equivalent in the sense that they achieve the exact same value f^* of the objective.*

Theorem 1 establishes that for each initial point (C^0, D^0), the bounded iterate sequence in Algorithm 1 is such that all its accumulation points achieve the same value f^* of the objective. They are equivalent in that sense. The value of f^* could vary with different initializations (C^0, D^0). We thus have the following corollary of Theorem 1 that holds because the distance between a bounded sequence and its (non-empty and compact) set of accumulation points converges to zero.

Corollary 1 *For each (C^0, D^0), the iterate sequence in Algorithm 1 converges to an equivalence class of accumulation points.*

The following Theorem 2 considers a very special case of SOUP dictionary learning, where the dictionary has a single atom. In this case, the SOUP learning Problem (P1) is the problem of obtaining a sparse rank-one approximation of the training matrix Y. In this case, Theorem 2 establishes that the iterates in Algorithm 1 converge to the set of critical points (i.e., the distance between the iterates and the set converges to zero) of the objective f.

Theorem 2 *Consider Algorithm 1 with $J = 1$. Let $\{c^t, d^t\}$ denote the bounded iterate sequence generated by Algorithm 1 in this case with data $Y \in \mathbb{R}^{n \times N}$ and initial (c^0, d^0). Then, every accumulation point of the iterate sequence is a critical point of the objective f, i.e., the iterates converge to the set of critical points of f.*

For the general case ($J \neq 1$), we have the following results.

Theorem 3 *Let $\{C^t, D^t\}$ denote the bounded iterate sequence generated by Algorithm 1 with training data $Y \in \mathbb{R}^{n \times N}$ and initial (C^0, D^0). Suppose each accumulation point (C, D) of the iterate sequence is such that the matrix B with columns $b_j = E_j^T d_j$ and $E_j = Y - DC^T + d_j c_j^T$, has no entry with magnitude λ. Then every accumulation point of the iterate sequence is a critical point of the objective $f(C, D)$. Moreover, the two sequences with terms $\left\|D^t - D^{t-1}\right\|_F$ and $\left\|C^t - C^{t-1}\right\|_F$ respectively, both converge to zero.*

Theorem 3 establishes that the iterates in SOUP-DIL converge to the set of critical points of $f(C, D)$. For each initial (C^0, D^0), the iterate sequence in Algorithm 1 converges (using Corollary 1) to an equivalence class of critical points of f. Theorem 3 also establishes that $\left\|D^t - D^{t-1}\right\|_F \to 0$ and

Barbara Boat Couple Hill Lena

Figure 2: The 512×512 standard images used in our experiments.

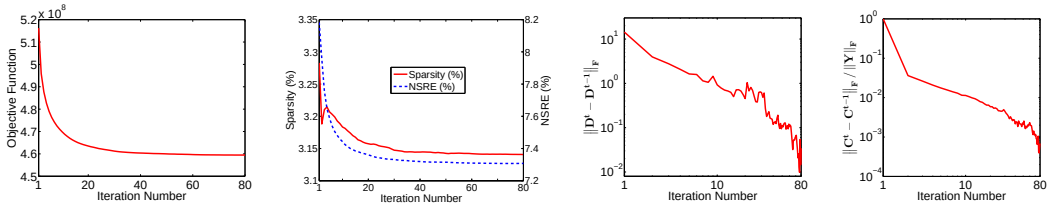

(a) (b) (c) (d)

Figure 3: Convergence of Algorithm 1: (a) Objective function; (b) NSRE (percentage) and sparsity factor of C (as a percentage); (c) changes between successive D iterates ($\left\|D^t - D^{t-1}\right\|_F$); and (d) normalized changes between successive C iterates ($\left\|C^t - C^{t-1}\right\|_F / \|Y\|_F$).

$\left\|C^t - C^{t-1}\right\|_F \to 0$. These are necessary but not sufficient conditions for the convergence of the entire sequences $\{D^t\}$ and $\{C^t\}$. The assumption on the entries of the matrix B in Theorem 3 is equivalent to assuming that for every $1 \leq j \leq J$, there is a unique minimizer of f with respect to c_j with all other variables fixed to their values in the accumulation point (C, D).

Although Theorem 3 uses a uniqueness condition, the following conjecture postulates that provided the following Assumption 1 (that uses a probabilistic model for the data) holds, the uniqueness condition holds with probability 1, i.e., the probability of a tie in assigning sparse codes is zero.

Assumption 1. The training signals $y_i \in \mathbb{R}^n$ for $1 \leq i \leq N$, are drawn independently from an absolutely continuous probability measure over the n-dimensional ball $S \triangleq \{y \in \mathbb{R}^n : \|y\|_2 \leq c_0\}$ for some $c_0 > 0$.

Conjecture 1 *Let Assumption 1 hold. Then, with probability 1, every accumulation point (C, D) of the iterate sequence in Algorithm 1 is such that for each $1 \leq j \leq J$, the minimizer of $f(c_1, ..., c_{j-1}, \tilde{c}_j, c_{j+1}, ..., c_J, d_1, ..., d_J)$ with respect to $\tilde{c}_j$ is unique.*

If Conjecture 1 holds, then every accumulation point of the iterate sequence in Algorithm 1 is immediately a critical point of $f(C, D)$ with probability 1.

4 Numerical Experiments

This section presents numerical results illustrating the practical convergence behavior of the proposed algorithm, as well as its usefulness in sparse signal representation and image denoising.

4.1 Convergence Experiment

To study the convergence behavior of the proposed SOUP-DIL Algorithm 1, we extracted 3×10^4 patches of size 8×8 from randomly chosen locations in the images Barbara, Boat, and Hill, shown in Fig. 2. The SOUP-DIL Algorithm 1 for (P1) was then used to learn a 64×256 overcomplete dictionary for this data, with $\lambda = 69$. Algorithm 1 is initialized as mentioned in Section 2.3.

Fig. 3 illustrates the convergence behavior of SOUP-DIL. The objective in our method (Fig. 3(a)) converged monotonically and quickly over the iterations. We define the *normalized sparse representation error* (NSRE) as $\left\|Y - DC^T\right\|_F / \|Y\|_F$, which is used to measure the sparse representation

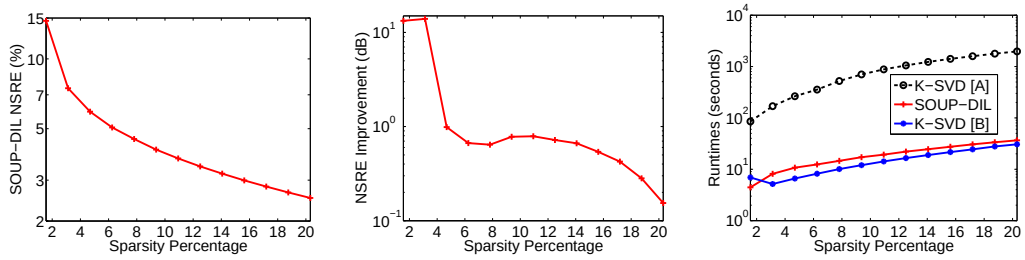

(a) (b) (c)

Figure 4: Behavior of SOUP-DIL Algorithm 1 and K-SVD at various sparsities: (a) NSRE (percentage) for SOUP-DIL; (b) Improvements in NSRE (in decibels) provided by SOUP-DIL over K-SVD; and (c) runtimes for SOUP-DIL, the matlab implementation of K-SVD (Elad, 2009) (labeled [A]), and the efficient (partial) MEX/C implementation of K-SVD (Rubinstein, 2010) (labeled [B]).

performance of the learned dictonaries. Fig. 3(b) shows the normalized sparse representation error and sparsity factor (for C), both expressed as percentages, for Algorithm 1. Both these components of the objective converged quickly, and the NSRE improved by 1 dB (decibel) beyond the first iteration, indicating the success of the SOUP-DIL Algorithm 1 in representing data using a small number of non-zero coefficients (sparsity factor of 3.14% at convergence).

Importantly, both the quantities $\left\|D^t - D^{t-1}\right\|_F$ (Fig. 3(c)) and $\left\|C^t - C^{t-1}\right\|_F$ (Fig. 3(d)) converge towards 0, as predicted by Theorem 3. This implies that $\left\|Z^t - Z^{t-1}\right\|_F$ with $Z^t = D^t \left(C^t\right)^T$ converges towards zero too. The above results are indicative (are necessary but not sufficient conditions) of the convergence of the entire sequences $\{D^t\}$, $\{C^t\}$, and $\{Z^t\}$ for our algorithm in practice. In contrast, Bao et al. (2014) showed that the distance between successive iterates may not converge to 0 for popular algorithms such as K-SVD.

4.2 Sparse Representation of Data

The second experiment worked with the same data as in Section 4.1 and learned dictionaries of size 64×256 for various choices of the parameter λ in (P1). We compare the performance of our algorithm to the K-SVD dictionary learning scheme (Aharon et al., 2006). The K-SVD method was executed for several choices of sparsity (number of non-zeros) of the columns of C^T. The λ values for (P1) were chosen so as to achieve similar average column sparsity levels in C^T as K-SVD. Both algorithms were initialized with the same overcomplete DCT and ran for 10 iterations each.

Fig. 4 shows the behavior of the SOUP-DIL and K-SVD algorithms for average column sparsity levels of C^T ranging from about 2% to 20%. As expected, the NSRE values for the SOUP-DIL algorithm decreased monotonically (Fig. 4(a)) when the average column sparsity levels in C^T increased (i.e., as λ decreased). Importantly, the SOUP-DIL algorithm provided better data representations (Fig. 4(b)) than K-SVD at the various tested sparsity levels. Large improvements of about 14 dB and 1 dB are observed at low and mid sparsity levels, respectively.

Fig. 4(c) compares the runtimes of the unoptimized Matlab implementation of our method to those of the unoptimized Matlab implementation of K-SVD (Elad, 2009) as well as the efficient (partial) MEX/C implementation (Rubinstein, 2010; Rubinstein et al., 2008) of K-SVD demonstrating large speedups of 40-50 times for SOUP-DIL over the first K-SVD implementation (at most sparsities), while the runtimes for the SOUP-DIL method are about the same as those of the second K-SVD implementation. Since these results were obtained using only an unoptimized Matlab implementation of SOUP-DIL, we expect significant speed-ups for our scheme with code optimization or C/C++ implementations. All computations in this work were performed with an Intel Xeon CPU X3230 at 2.66 GHz and 8 GB memory, employing a 64-bit Windows 7 operating system.

4.3 Image Denoising

In image denoising, the goal is to recover an estimate of an image $x \in \mathbb{R}^M$ (2D image represented as a vector) from its corrupted measurements $y = x + h$, where h is the noise. For example, the entries of h may be i.i.d. Gaussian with zero mean and standard deviation σ. To perform image denoising

using (P1), we first extract all the overlapping patches (with maximum overlap) of the noisy image y, and construct the training set $Y \in \mathbb{R}^{n \times N}$ as a matrix whose columns are those noisy patches. We then use Problem (P1) to learn a dictionary and sparse codes for Y, with the weight $\lambda \propto \sigma$. To obtain the denoised image estimate, we then solve the following least squares problem, where $\hat{D}$ and $\hat{\alpha}_j$ denote the learned dictionary and patch sparse codes obtained from the noisy patches, and P_j is an operator that extracts a patch as a vector:

$$\min_x \sum_{j=1}^{N} \left\| P_j x - \hat{D}\hat{\alpha}_j \right\|_2^2 + \nu \left\| x - y \right\|_2^2 \tag{10}$$

The optimal $\hat{x}$ in (10) is easily obtained by summing together the denoised patch estimates $\hat{D}\hat{\alpha}_j$ at their 2D locations, and computing a weighted average between this result and the noisy image.

K-SVD based denoising[3] (Elad & Aharon, 2006) involves a similar methodology as described above for (P1), but differs in the dictionary learning procedure, where the ℓ_0 "norms" of the sparse codes are minimized so that a fitting constraint or error constraint of $\left\| P_j y - \hat{D}\hat{\alpha}_j \right\|_2^2 \leq \epsilon$ (e.g., $\epsilon = nC^2\sigma^2$ for i.i.d. Gaussian noise, with $C > 1$ a constant) is met for representing the noisy patches. This constraint serves as a strong prior, and is a key reason for the denoising capability of K-SVD. Hence, in our method based on (P1), once the dictionary $\hat{D}$ is learned from noisy patches, we estimate the patch sparse codes $\hat{\alpha}_j$ using a single pass (over the noisy patches) of orthogonal matching pursuit (OMP) by employing an error constraint criterion like in K-SVD. Although we only use information on the noise statistics in a sub-optimal way for (P1), we still show good denoising performance (vis-a-vis K-SVD) in the following with this approach.

For the denoising experiments, we work with the images in Fig. 2, and simulate i.i.d. Gaussian noise at six different noise levels ($\sigma = 5$, 10, 20, 25, 30, 100) for each of the images. We then denoise these images using both the SOUP-DIL denoising method outlined above and K-SVD. SOUP-DIL denoising uses 8×8 patches, a 64×256 dictionary, $\lambda = 5\sigma$, $\nu = 20/\sigma$ (in (10)), an overcomplete DCT initial dictionary in learning, and 10 iterations of Algorithm 1. These settings were found to work well for our method.

Table 1 lists the denoising PSNRs obtained by the SOUP-DIL denoising method, along with the PSNRs obtained by K-SVD. For comparison, we also list the denoising PSNRs obtained by employing the overcomplete DCT dictionary for denoising. In the latter case, the same strategy (and parameter settings) as used by K-SVD based denoising is adopted but while skipping the learning process.

The various methods denoise about the same in Table 1 for a low noise level of $\sigma = 5$. For $\sigma > 5$, the SOUP-DIL method denoises about 0.4 dB better on the average than the overcomplete DCT. While the SOUP-DIL method and K-SVD denoise about the same at low and mid noise levels, the SOUP-DIL scheme performs about 0.1 dB better on average at higher noise levels such as $\sigma = 30$ or 100. Importantly, SOUP-DIL based denoising is efficient and the learning procedure has good convergence properties.

Table 2 compares the average runtimes (averaged over the images) of the (unoptimized) Matlab implementations of K-SVD denoising (Elad, 2009) and SOUP-DIL denoising at various noise levels. The runtimes for the dictionary learning step and for denoising with the learned dictionary are individually shown. The latter step is similar for both K-SVD and SOUP-DIL, and involves sparse coding (with OMP) all image patches using the learned dictionary and solving (10). SOUP-DIL clearly achieves low runtimes for the dictionary learning step, while it has similar runtimes as K-SVD for the step of denoising with the learned dictionary.

Fig. 5 shows the dictionaries learned from noisy patches using the SOUP-DIL method and K-SVD for the image Barbara at $\sigma = 20$. Both dictionaries show frequency and textural features that are

[3]The K-SVD method is a highly popular dictionary learning scheme that has been applied to a wide variety of image processing applications including denoising (Elad & Aharon, 2006). Mairal et al. (2009) proposed a non-local method for image denoising that also exploits learned dictionaries and achieves denoising performance comparable to the well-known BM3D (Dabov et al., 2007) denoising method. Similar extensions to our proposed method in denoising and other applications is beyond the scope of this work and is left for future work. For K-SVD denoising, we used the built-in parameter settings of the author's implementation (Elad, 2009).

Img.	σ	Noisy	O-DCT	K-SVD	SOUP
1	5	34.16	37.25	**37.29**	37.28
	10	28.11	33.40	33.49	**33.50**
	20	22.11	29.71	**30.01**	29.99
	25	20.17	28.53	28.88	**28.92**
	30	18.58	27.53	27.88	**27.97**
	100	8.13	22.59	22.58	**22.71**
2	5	34.15	37.94	**38.08**	38.04
	10	28.14	33.96	**34.43**	34.37
	20	22.13	29.95	**30.83**	30.79
	25	20.17	28.68	29.63	**29.64**
	30	18.59	27.62	28.54	**28.63**
	100	8.11	21.87	21.87	**21.97**
3	5	34.15	37.09	**37.21**	37.16
	10	28.13	33.43	**33.62**	33.60
	20	22.10	29.92	30.36	**30.37**
	25	20.17	28.79	29.28	**29.30**
	30	18.60	27.93	28.41	**28.43**
	100	8.13	22.79	22.81	**22.96**

Img.	σ	Noisy	O-DCT	K-SVD	SOUP
4	5	34.15	37.02	**37.08**	37.05
	10	28.14	33.26	**33.45**	33.44
	20	22.10	29.85	30.17	**30.20**
	25	20.18	28.89	29.23	**29.31**
	30	18.57	28.14	28.43	**28.56**
	100	8.16	24.00	23.98	**24.03**
5	5	34.16	38.52	**38.62**	38.55
	10	28.12	35.30	**35.48**	35.47
	20	22.11	32.02	32.40	**32.40**
	25	20.18	30.89	31.32	**31.32**
	30	18.59	29.98	30.41	**30.46**
	100	8.14	24.45	24.51	**24.63**
Avg.	5	34.16	37.56	**37.66**	37.61
	10	28.13	33.87	**34.09**	34.07
	20	22.11	30.29	30.75	**30.75**
	25	20.17	29.16	29.67	**29.70**
	30	18.58	28.24	28.74	**28.81**
	100	8.13	23.14	23.15	**23.26**

Table 1: PSNR values in decibels for denoising with overcomplete DCT (O-DCT), K-SVD, and SOUP-DIL. All dictionaries have size 64×256. The PSNR values of the noisy images (Noisy) are also shown. The images 1 to 5 correspond to Couple, Barbara, Boat, Hill, and Lena, respectively. The last block of entries in the table corresponds to the average (over the images) PSNR values. The best PSNR in each row is marked in bold.

Method	Step	$\sigma = 5$	$\sigma = 10$	$\sigma = 20$	$\sigma = 25$	$\sigma = 30$	$\sigma = 100$
K-SVD	Dictionary Learning	1587	828	198	144	113	45
	Denoising with learned D	632	171	56	42	34	15
	Total	2219	999	254	186	147	60
SOUP-DIL	Dictionary Learning	132	71	46	42	40	33
	Denoising with learned D	438	159	62	51	39	17
	Total	**570**	**230**	**108**	**93**	**79**	**50**

Table 2: Runtimes (seconds) averaged over the images at each noise level for the Matlab implementations of K-SVD (Elad, 2009) and SOUP-DIL denoising. The runtimes are broken into the time for the dictionary learning step, and the time for denoising with the learned dictionary. The latter step is similar for both K-SVD and SOUP-DIL, and involves sparse coding (with OMP) all image patches using the learned dictionary and solving (10). The lowest total times are bolded.

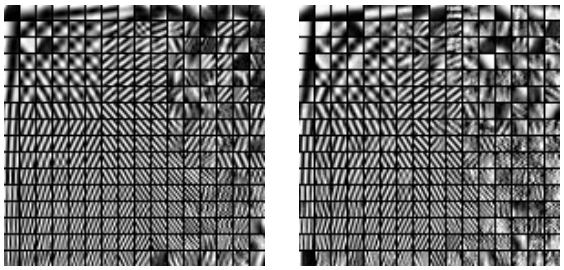

(a) (b)

Figure 5: Learned dictionaries for Barbara at $\sigma = 20$: (a) SOUP dictionary; and (b) K-SVD dictionary. The columns of the dictionaries are shown as patches.

specific to the image Barbara. By learning such image-specific features, the SOUP-DIL method (and K-SVD) easily outperforms fixed dictionaries such as the overcomplete DCT in denoising.

5 CONCLUSIONS

This paper proposed a fast block coordinate descent method for synthesis dictionary learning with an ℓ_0 penalty. The key idea is to represent the training data as a sum of sparse rank-one matrices, and then efficiently estimate the factors of the rank-one matrices. A convergence analysis was presented for the proposed block coordinate descent algorithm for a highly non-convex problem. The proposed approach had comparable or superior performance and significant speed-ups over the classical K-

SVD method in sparse signal representation and denoising. Extensions of the method for online learning (Mairal et al., 2010) merit further study.

ACKNOWLEDGMENTS

This work was supported in part by the following grants: ONR grant N00014-15-1-2141, DARPA Young Faculty Award D14AP00086, ARO MURI grants W911NF-11-1-0391 and 2015-05174-05, NIH grant U01 EB01875301, and a UM-SJTU seed grant.

REFERENCES

Agarwal, A., Anandkumar, A., Jain, P., Netrapalli, P., and Tandon, R. Learning sparsely used overcomplete dictionaries. *Journal of Machine Learning Research*, 35:1–15, 2014.

Aharon, M., Elad, M., and Bruckstein, A. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing*, 54(11):4311–4322, 2006.

Arora, S., Ge, R., and Moitra, A. New algorithms for learning incoherent and overcomplete dictionaries. 2013. Preprint: `http://arxiv.org/pdf/1308.6273v5.pdf`.

Bao, C., Ji, H., Quan, Y., and Shen, Z. L0 norm based dictionary learning by proximal methods with global convergence. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3858–3865, 2014.

Bao, C., Ji, H., Quan, Y., and Shen, Z. Dictionary learning for sparse coding: Algorithms and analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2015.

Chen, S. S., Donoho, D. L., and Saunders, M. A. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.*, 20(1):33–61, 1998.

Dabov, K., Foi, A., Katkovnik, V., and Egiazarian, K. Image denoising by sparse 3D transform-domain collaborative filtering. *IEEE Trans. on Image Processing*, 16(8):2080–2095, 2007.

Elad, M. and Aharon, M. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Trans. Image Process.*, 15(12):3736–3745, 2006.

Elad, Michael. `http://www.cs.technion.ac.il/~elad/software/`, 2009. [Online; accessed Nov. 2015].

Engan, K., Aase, S.O., and Hakon-Husoy, J.H. Method of optimal directions for frame design. In *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 2443–2446, 1999.

Grama, A., Gupta, A., Karypis, G., and Kumar, V. *Introduction to Parallel Computing*. Addison Wesley, Harlow, England, 2003.

Gribonval, R. and Schnass, K. Dictionary identification–sparse matrix-factorization via l_1 -minimization. *IEEE Trans. Inform. Theory*, 56(7):3523–3539, 2010.

Mairal, J., Bach, F., Ponce, J., Sapiro, G., and Zisserman, A. Non-local sparse models for image restoration. In *IEEE International Conference on Computer Vision*, pp. 2272–2279, Sept 2009.

Mairal, J., Bach, F., Ponce, J., and Sapiro, G. Online learning for matrix factorization and sparse coding. *J. Mach. Learn. Res.*, 11:19–60, 2010.

Mordukhovich, B. S. *Variational Analysis and Generalized Differentiation. Vol. I: Basic theory*. Springer-Verlag, 2006.

Needell, D. and Tropp, J.A. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26(3):301 – 321, 2009.

Pati, Y., Rezaiifar, R., and Krishnaprasad, P. Orthogonal matching pursuit : recursive function approximation with applications to wavelet decomposition. In *Asilomar Conf. on Signals, Systems and Comput.*, pp. 40–44 vol.1, 1993.

Rakotomamonjy, A. Direct optimization of the dictionary learning problem. *IEEE Transactions on Signal Processing*, 61(22):5495–5506, 2013.

Ravishankar, S. and Bresler, Y. MR image reconstruction from highly undersampled k-space data by dictionary learning. *IEEE Trans. Med. Imag.*, 30(5):1028–1041, 2011.

Ravishankar, S. and Bresler, Y. Efficient blind compressed sensing using sparsifying transforms with convergence guarantees and application to magnetic resonance imaging. *SIAM Journal on Imaging Sciences*, 8(4):2519–2557, 2015.

Ravishankar, S., Nadakuditi, R. R., and Fessler, J. A. Efficient sum of outer products dictionary learning (SOUP-DIL) - the ℓ_0 method. 2015. Preprint: `http://arxiv.org/abs/1511.08842`.

Rockafellar, R. T. and Wets, Roger J.-B. *Variational Analysis, 1st edn.* Springer-Verlag, 1997.

Rubinstein, R., Zibulevsky, M., and Elad, M. Efficient implementation of the k-svd algorithm using batch orthogonal matching pursuit. `http://www.cs.technion.ac.il/~ronrubin/Publications/KSVD-OMP-v2.pdf`, 2008. Technion - Computer Science Department - Technical Report.

Rubinstein, R., Zibulevsky, M., and Elad, M. Double sparsity: Learning sparse dictionaries for sparse signal approximation. *IEEE Transactions on Signal Processing*, 58(3):1553–1564, 2010.

Rubinstein, Ron. `http://www.cs.technion.ac.il/~ronrubin/software.html`, 2010. [Online; accessed Nov. 2015].

Sadeghi, M., Babaie-Zadeh, M., and Jutten, C. Dictionary learning for sparse representation: A novel approach. *IEEE Signal Processing Letters*, 20(12):1195–1198, Dec 2013.

Sadeghi, M., Babaie-Zadeh, M., and Jutten, C. Learning overcomplete dictionaries based on atom-by-atom updating. *IEEE Transactions on Signal Processing*, 62(4):883–891, 2014.

Skretting, K. and Engan, K. Recursive least squares dictionary learning algorithm. *IEEE Transactions on Signal Processing*, 58(4):2121–2130, 2010.

Smith, L. N. and Elad, M. Improving dictionary learning: Multiple dictionary updates and coefficient reuse. *IEEE Signal Processing Letters*, 20(1):79–82, Jan 2013.

Spielman, D. A., Wang, H., and Wright, J. Exact recovery of sparsely-used dictionaries. In *Proceedings of the 25th Annual Conference on Learning Theory*, pp. 37.1–37.18, 2012.

Tseng, P. Convergence of a block coordinate descent method for nondifferentiable minimization. *J. Optim. Theory Appl.*, 109(3):475–494, 2001.

Xu, Y. and Yin, W. A fast patch-dictionary method for whole-image recovery. 2013. URL `ftp://ftp.math.ucla.edu/pub/camreport/cam13-38.pdf`. UCLA CAM report 13-38.

Yaghoobi, M., Blumensath, T., and Davies, M. Dictionary learning for sparse approximations with the majorization method. *IEEE Transaction on Signal Processing*, 57(6):2178–2191, 2009.

A PROOF OF PROPOSITION 1

First, for a vector d_j that has unit ℓ_2 norm, we have the following equality

$$\begin{aligned}\left\|E_j - d_j c_j^T\right\|_F^2 &= \|E_j\|_F^2 + \|c_j\|_2^2 - 2c_j^T E_j^T d_j \\ &= \left\|c_j - E_j^T d_j\right\|_2^2 + \|E_j\|_F^2 - \left\|E_j^T d_j\right\|_F^2 \end{aligned} \tag{11}$$

By substituting (11) into (1), it is clear that Problem (1) is equivalent to

$$\min_{c_j} \left\|c_j - E_j^T d_j\right\|_2^2 + \lambda^2 \|c_j\|_0 \;\; \text{s.t.} \;\; \|c_j\|_\infty \leq L. \tag{12}$$

Define $b \triangleq E_j^T d_j$. Then, the objective in (12) simplifies to $\sum_{i=1}^N \left\{|c_{ji} - b_i|^2 + \lambda^2 \theta(c_{ji})\right\}$ with

$$\theta(a) = \begin{cases} 0, & \text{if } a = 0 \\ 1, & \text{if } a \neq 0 \end{cases} \tag{13}$$

Therefore, we solve for each entry c_{ji} of c_j as

$$\hat{c}_{ji} = \arg\min_{c_{ji}} |c_{ji} - b_i|^2 + \lambda^2 \theta(c_{ji}) \;\; \text{s.t.} \;\; |c_{ji}| \leq L. \tag{14}$$

It is straightforward to show that when $|b_i| \leq L$ (case (a)),

$$\hat{c}_{ji} = \begin{cases} 0, & \text{if } b_i^2 < \lambda^2 \\ b_i, & \text{if } b_i^2 > \lambda^2 \end{cases} \tag{15}$$

When $|b_i| = \lambda$ ($\lambda < L$), the optimal $\hat{c}_{ji}$ can be either b_i or 0 (non-unique), and both these settings achieve the minimum objective value λ^2 in (14). Next, when $|b_i| > L$ (case (b)), we have

$$\hat{c}_{ji} = \begin{cases} 0, & \text{if } b_i^2 < \left(L\,\text{sign}(b_i) - b_i\right)^2 + \lambda^2 \\ L\,\text{sign}(b_i), & \text{if } b_i^2 > \left(L\,\text{sign}(b_i) - b_i\right)^2 + \lambda^2 \end{cases} \tag{16}$$

Since $L > \lambda$, clearly $b_i^2 > \left(L\,\text{sign}(b_i) - b_i\right)^2 + \lambda^2$ in (16).

Thus, an optimal $\hat{c}_{ji}$ in (14) is compactly written as $\hat{c}_{ji} = \min\left(\left|H_\lambda(b_i)\right|, L\right) \cdot \text{sign}(b_i)$, thereby establishing (3). The condition for uniqueness of the sparse coding solution follows from the arguments for case (a) above. ■

B PROOF OF PROPOSITION 2

First, for a vector d_j that has unit ℓ_2 norm, the following holds:

$$\left\|E_j - d_j c_j^T\right\|_F^2 = \|E_j\|_F^2 + \|c_j\|_2^2 - 2\, d_j^T E_j c_j \tag{17}$$

Upon substituting (17) into (4), Problem (4) simplifies to

$$\max_{d_j} \; d_j^T E_j c_j \;\; \text{s.t.} \;\; \|d_j\|_2 = 1. \tag{18}$$

By the Cauchy Schwarz inequality $d_j^T E_j c_j \leq \|E_j c_j\|_2$ for unit norm d_j. Thus, a solution to (18) that achieves the value $\|E_j c_j\|_2$ for the objective is

$$\hat{d}_j = \begin{cases} \frac{E_j c_j}{\|E_j c_j\|_2}, & \text{if } E_j c_j \neq 0 \\ v_1, & \text{if } E_j c_j = 0 \end{cases} \tag{19}$$

Obviously, any $d \in \mathbb{R}^n$ would be a minimizer (non-unique) in (18) when $E_j c_j = 0$. In particular $\hat{d}_j = v_1$ works.

Next, we show that $E_j c_j = 0$ in our algorithm if and only if $c_j = 0$. This result together with (19) immediately establishes the proposition. Since the c_j used in the dictionary atom update step (4) was obtained as a minimizer in the preceding sparse coding step (1), we have the following inequality for all $c \in \mathbb{R}^N$ with $\|c\|_\infty \leq L$, and $\tilde{d}_j$ denotes the jth atom in the preceding sparse coding step:

$$\left\|E_j - \tilde{d}_j c_j^T\right\|_F^2 + \lambda^2 \|c_j\|_0 \leq \left\|E_j - \tilde{d}_j c^T\right\|_F^2 + \lambda^2 \|c\|_0 \tag{20}$$

If $E_j c_j = 0$, the left hand side above simplifies to $\|E_j\|_F^2 + \|c_j\|_2^2 + \lambda^2 \|c_j\|_0$, which is clearly minimal when $c_j = 0$. Thus, when $E_j c_j = 0$, we must also have $c_j = 0$. ■